Learning and Tuning Meta-heuristics in Plan Space Planning

Shashank Shekhar and Deepak Khemani
Department of Computer Science & Engineering
IIT Madras, Chennai
Tamil Nadu, India 600 036
{sshekhar,khemani}@cse.iitm.ac.in

arXiv:1601.07483v3 [cs.AI] 24 Apr 2016

Abstract

In recent years, the planning community has observed that techniques for learning heuristic functions have yielded improvements in performance. One approach is to use offline learning to learn predictive models from existing heuristics in a domain dependent manner. These learned models are deployed as new heuristic functions. The learned models can in turn be tuned online using a domain independent error correction approach to further enhance their informativeness. The online tuning approach is domain independent but instance specific, and contributes to improved performance for individual instances as planning proceeds. Consequently it is more effective in larger problems.

In this paper, we mention two approaches applicable in Partial Order Causal Link (POCL) Planning that is also known as Plan Space Planning. First, we endeavour to enhance the performance of a POCL planner by giving an algorithm for supervised learning. Second, we then discuss an online error minimization approach in POCL framework to minimize the step-error associated with the offline learned models thus enhancing their informativeness. Our evaluation shows that the learning approaches scale up the performance of the planner over standard benchmarks, specially for larger problems.

Introduction

In the recent International Planning Competitions (IPC) state-space based and total-order planners like LAMA (Richter and Westphal 2010), Fast Downward Stone Soup (Helmert, Röger, and Karpas 2011), and Fast Downward Stone Soup 2014 (Röger, Pommerening, and Seipp 2014) have performed well. These planners are very efficient, generate consistent states fast, and use powerful state-based heuristics. But they often commit too early to ordering of actions, giving up on flexibility in the generated plans. In contrast, the POCL framework (Coles et al. 2010) generates more flexible plans, but in general is computationally more intensive than the state-space based approaches. The POCL framework has found applicability in multi-agent planning (Kvarnström 2011) and temporal planning (Benton, Coles, and Coles 2012).

Researchers have recently investigated the applicability of state-space heuristic learning approaches (Sapena, Onaindıa, and Torreno 2014; Coles et al. 2010) in POCL planning. This revival of interest is due to the idea of delayed commitment of RePOP (Nguyen and Kambhampati 2001) and VHPOP (Younes and Simmons 2003). In this paper we further investigate the adaptation of state space approaches in POCL planners yielding quality plans over the same or even larger problems. In general, due to the diverse nature of planning problems characterized by the degree of interactions between subgoals, a heuristic function does not always work well in all planning domains. Various techniques have been devised to increase the informativeness of heuristics in the state space arena. One approach strengthens the delete relaxation heuristic by incrementing lower bounds to tighten the admissibility of the heuristic, repeatedly by solving relaxed versions of a planning problem (Haslum 2012). In another approach, to circumvent the trade-offs of combining multiple heuristics, a decision rule is used for selecting a heuristic function in a given state. An active online learning technique is applied to learn a model for that given decision rule (Domshlak, Karpas, and Markovitch 2010).

We present a couple of machine learning techniques, influenced from (Arfaee, Zilles, and Holte 2011; Thayer, Dionne, and Ruml 2011; Samadi, Felner, and Schaeffer 2008; Virseda, Borrajo, and Alcázar 2013), which improve the effectiveness of heuristics in the POCL framework. First, we apply domain wise regression techniques in a supervised manner, using existing POCL heuristics as features. For target generation we use our planner named RegPOCL. This is based on VHPOP and uses grounded actions to speed up the planning process. Next, we further give another technique for POCL framework which enhances the informativeness of these offline learned models. This technique is an adapted version of an online heuristic adjustment approach based on temporal difference (TD) learning (Sutton 1988; Thayer, Dionne, and Ruml 2011). We extend this domain independent approach for learning instance specific details, which corrects the error associated with the learned models, thus making them more informed. The RegPOCL planner employs these two approaches and evaluation shows that it is more efficient on the benchmarks. We have confined the evaluation to non-temporal STRIPS domains.

The rest of the paper is structured as follows. After looking at the motivation for exploring the POCL approach, we describe the learning approaches used along with the experi-

Copyright © 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

ment design. These are followed by the experimental results and the concluding discussion.

POCL Planning

A POCL planner starts search with a null partial plan and progresses over a space of partial plans, by adding a resolver to a partial plan to remove a flaw. We use heuristics from (Younes and Simmons 2003) for selecting the most adverse flaw to resolve in the selected partial plan. The POCL framework has certain advantages over forward state-space search (FSSS). FSSS has the problem of premature commitment to an ordering between two actions which reduces the plan flexibility. It does so to avoid any mutual interference between actions, though there may not be any. POCL planning avoids committing unnecessarily to ordering actions.

FSSS faces problems while solving instances with deadlines. Deadlines may arise within the interval(s) of one or more durative actions. In general, the actions may produce some delayed effects, and this may have ramifications on deadlines as well, which creates deadlines relative to their starting points (Coles et al. 2010). FSSS also suffers from significant backtracking as it may explore all possible plan permutations in the absence of effective guidance. The POCL framework also has several advantages in temporal planning, specially in planning with complex temporal constraints beyond actions with duration. These limitations of FSSS motivate us to explore the POCL framework.

Example. Suppose we are required to add four actions [a_1, a_2, a_3, a_4] to a plan, where a_2 is dependent on a_1 and a_4 is dependent on a_3. There is no interference between any two actions apart from the above dependencies. In this case, FSSS gives an ordering or timestamp [0, 1, 2, 3], with a makespan 4, whilst the delayed commitment strategy would give more choices with flexibility in the orderings like [2, 0, 1, 3] and [0, 2, 1, 3]. If parallel execution is allowed, makespan would be 2. If another action a_5, which is dependent on a_3, has to be introduced in the plan then FSSS will allot it a timestamp 4, whereas delayed-commitment strategy could allot it 1.

However, if we ignore the absence of the flexibility and action parallelism in FSSS, it is very fast in recovering from a situation that would arise due to committing to some wrong choices during planning. FSSS has the advantage of faster search state generation and powerful state-based heuristics.

Learning Approaches Used

We propose a two fold approach to learn better heuristic functions. First, existing heuristic function are combined by a process of offline learning that generates learned predictive models. This is followed by an online technique of adjusting the step-error associated with these models during partial plan refinement. We divide this section into two parts: the first describes the offline learning techniques to perform regression, and the second the technique of further strengthening the informativeness of a learned predictive model.

Offline Learning

Offline learning is an attractive approach because generating training sets in most planning domains is a fast and simple process. The learning process has three phases: (*i*) dataset preparation, (*ii*) training, and (*iii*) testing. The training instances gathered by solving small problems become the inputs to the used regression techniques (*e.g.* linear regression and M5P, described later), which learn effective ways of combining existing heuristic functions. The testing phase is used to validate the best ways of combining the known heuristic functions. Algorithm 1, described below, embodies the complete training process.

Algorithm 1 The algorithm used during training phase

```
1:  Input
2:      AS - Attribute Selection; T - Training Dataset;
3:      S - Problem Set; L - Learning Technique;
4:      H - Heuristic Set; RegPOCL - The Planner.
5:  Output
6:      M : T → ℝ   // A learned predictive model
7:  procedure TRAINING-ALGORITHM(AS, T, L)
8:      T ← φ   // Domain specific.
9:      T ← DATASET-PREP(RegPOCL, S, H, T)
10:     Training Instances ← Apply(AS, T)
11:     return M ← Apply(L, Training Instances)
12: end procedure

13: procedure DATASET-PREP(RegPOCL, S, H, T)
14:     F - Check; 𝒯 - Target Value;
15:     Π - A set of seed partial plans.
16:     for each p ∈ S do
17:         Π ← Null partial plan of p
18:         for a random sp ∈ Π do
19:             (F, 𝒯, Π_loc) ← Solve(RegPOCL, sp, H)
20:             if F then   // sp refines completely.
21:                 Π ← Π ∪ Π_loc
22:                 Ins ← Comp-inst(RegPOCL, H, sp, 𝒯)
23:                 T ← T ∪ Ins
24:             end if
25:         end for   // Bounded number of iterations.
26:     end for
27:     return T
28: end procedure
```

Dataset Preparation The procedure DATASET-PREP() Line 13 in Algorithm 1 is used to solve a set (S) of planning problems. We consider only small problems that are gathered from each planning domain selected from previous IPCs. The output of Algorithm 1 is a trained predictive model (as shown in line 6). We consider each problem from S for the dataset preparation in each domain (line 16). In this algorithm, a seed partial plan is a new partial plan that gets generated due to a possible refinement of the selected partial plan. We select a seed sp from a set of seed partial plans Π (line 18). sp will be provided to RegPOCL for its further refinements. If RegPOCL is able to generate one consistent solution by refining sp completely using *Solve()* function (line 19), then the flag F will be true. We capture the

newly generated partial plans in local instance of Π called Π_{loc}. The target $\mathcal{T}$ captures the number of new actions that get added in sp. $\mathcal{T}$ is calculated when the planner refines sp completely using heuristics from H. Note that, $\mathcal{T}$ is not a heuristic value but the number of new actions added during the refinement process. The value of $\mathcal{T}$ is also the plan length found which might not be optimal. Since sp is refined completely, Π_{loc} is updated to Π (line 21). Line 22 computes a training instance Ins, using *Comp-inst()* function. For a given sp, the planner generates a training instance of the form $t(sp) = \langle\langle h_1(sp), h_2(sp), h_3(sp), h_4(sp), h_5(sp), h_6(sp)\rangle, \mathcal{T}\rangle$, where h_1 to h_6 are the feature heuristics. To maintain consistency, we update the training set T (line 23) only when RegPOCL refines the current seed sp completely. If complete refinement was not possible, all new seeds from Π_{loc} are dropped, even though it might be the case that Π_{loc} contains some consistent seeds particularly in the case of time out. To maintain diversity in T, for a given domain we randomly select a fixed number of seeds for the complete refinement process (line 18).

We execute Algorithm 1 once for each selected domain with a given set of feature heuristics. Note that learning does not guarantee optimal predictive models even though optimal targets have been used in the training (Virseda, Borrajo, and Alcázar 2013). Algorithm 1 hunts for a well informed heuristic using learning and does not bother about its admissibility. Since the state-of-the-art POCL heuristics are not optimal in nature (Younes and Simmons 2003), the usage of RegPOCL for generating training instances should not affect the performance of RegPOCL on large problems in the testing phase. The selection of RegPOCL for generating training sets might deteriorate the actual target values, as the targets calculated by RegPOCL are not optimal in general. Thus there is a possibility of learning inaccurate predictive models in the training phase, which might reduce the informativeness of the models. We enhance the informativeness of the models by correcting the step-error associated with them using an online heuristic tuning approach.

Training Once Algorithm 1 generates enough number of training instances for a given domain, it moves to the next step (line 10). We define a regression model $\mathrm{M} : \mathrm{T} \rightarrow \mathbb{R}$, where T is a training set and $\mathbb{R}$ is a set of real numbers. Following the general model training strategy, we use WEKA (Hall et al. 2009) to remove irrelevant or redundant attributes (line 10) from the training set. This reduces the effort of the planner because the planner must calculate the selected feature heuristics at each step of the planning process. Next, we apply model training process (line 11). We feed the training instances to different machine learning approaches to learn different predictive models.

Testing We test the predictive models on large problems. The models are directly compared to the current best heuristics in the POCL framework. For using machine learning approaches in planning efficiently, we select the best learned regression models and test the efficiency of RegPOCL by using them over different benchmarks. These models help in selecting the most suitable partial plan for refinement.

Offline learning learns a better model in terms of search guidance and accuracy than online learning (Samadi, Felner, and Schaeffer 2008; Thayer, Dionne, and Ruml 2011). An offline learned predictor is more accurate than an online one because in the offline case a complete training set is available. Another alternative to the above approaches would be bootstrapping methods (Arfaee, Zilles, and Holte 2011), where a set of problems is solved using a base heuristic within a specified time limit. Later, the solutions obtained for learning are used to generate a new more informed heuristic function.

Online Error Adjustment of A Predictive Model

The offline approach assumes that the learned knowledge can effectively be transferred from the training instances to the test instances. This is not always true as the planning problems are not always similar even though they belong to the same domain. Also, the training instances are required before the training process. For a planning domain, during the dataset preparation, calculation of features is not computationally hard, but calculating actual number of new actions $(\mathcal{T})$ needed for sp is relatively expensive. Online learning attempts to avoid these limitations. In our hybrid approach we use small instances for offline training, thus saving on time. This is followed up with online learning to improve the heuristic estimate on-the-fly. The online error tuning is based on temporal difference (TD) learning (Sutton 1988).

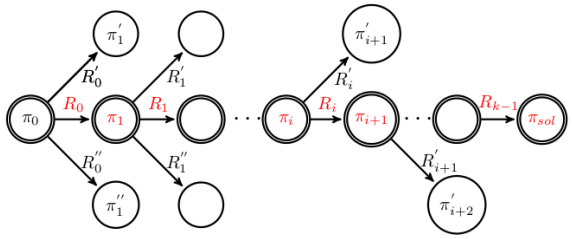

Figure 1: A POCL framework - The refinement *(R)* starts from π_0 and it goes to the solution plan (π_{sol}). At each step, for a selected partial plan, many refinements are possible like refinements of π_0 which lead to π_1, π_1', and π_1''. Here, the best child is shown in the horizontal refinements.

The online heuristic adjustment approach is inspired from a recent work presented as technical communication (Shekhar and Khemani 2015) tested in a few planning domains. We further develop this approach, and provide a complete derivation, and a proof (in Appendix) of the theorem used in the cited work. We assume that a predictive model (h) for a given partial plan (π) estimates the actual number of new actions required to be added in π to refine it completely to a solution plan, denoted by $h(\pi)$. Whilst $h^*(\pi)$ is the minimum number of new actions required for the same purpose (Nguyen and Kambhampati 2001). In the POCL framework, it is computationally expensive to predict $h^*(\pi)$. Since h is also likely to be uninformed, we do adjustments to h by observing single-step-error (ϵ_{h}) on-the-fly.

The minimum number of total refinements needed for a partial plan π_i to make it a solution plan, goes through its best child π_{i+1} that is obtained after refinement R_i. A child

π_{i+1} is the best child when it has the lowest prediction of the number of new actions needed for its complete refinement among its siblings (Figure 1). We break ties in favor of minimum number of actions in the children partial plans. The set of successors of a partial plan is potentially infinite. This is due to the introduction of loops in the plan which simply achieve and destroy subgoals (for example: $\langle$*(stack A B), (unstack A B)*$\rangle$ or $\langle$*(pickup A), (putdown A)*$\rangle$, in Blocksworld domain). Such looping is common during refinement process, specially when the heuristic is not well informed. We avoid such looping explicitly. This is crucial in real world scenarios. For example, a pair of actions like $\langle$*(load-truck obj truck loc), (unload-truck obj truck loc)*$\rangle$, in Driverlog domain, could be expensive. In general in plan space planning there is no backtracking. Each refinement of a partial plan leads to a different node in the plan space. This necessitates that we consider looping explicitly. Ideally,

$$h^*(\pi_i) = cost(R_i) + h^*(\pi_{i+1}) \tag{1}$$

In non-ideal cases where generating h^* is expensive, we assume that, $h^*(\pi_i) \approx h(\pi_i) \approx cost(R_i) + h(\pi_{i+1})$. In Figure 1, for a given partial plan π_i, suppose a minimum of n refinement steps are required to make π_i a solution plan. Here we assume that at each stage of the refinement process, the best child for each π_i is chosen, for all *i*. Here, n is an estimate of adding new actions in π_i. If we consider only unit cost actions then in the ideal case, $h^*(\pi_{i+1}) = h^*(\pi_i) - 1$. Generally, h commits some error at each refinement step called single-step-error (ϵ_h). The single-step-error associated with h, when π_i gets refined by an algorithm using h, denoted as $\epsilon_{h(\pi_i)}$ is,

$$\epsilon_{h(\pi_i)} = \big((1 + h(\pi_{i+1})) - h(\pi_i)\big) \tag{2}$$

Using Eq. (1), we derive (proof is in the Appendix) the following in a non-ideal case:

Theorem 1. *For a given learned predictive model (h) and partial plan (π_i) in Figure 1 which leads to the solution plan (π_{sol}) after certain refinement steps, the enhanced version of the predictive model (h^e) is,*

$$h^e(\pi_i) \;=\; h(\pi_i) \;+ \sum_{\pi' \textit{ from } \pi_i \rightsquigarrow \pi_{sol}} \epsilon_{h(\pi')} \tag{3}$$

where $\pi_i \rightsquigarrow \pi_{sol}$ is a path in Figure 1 which captures each partial plan (π') along the path between π_i and π_{sol}. The path includes π_i and excludes π_{sol}. The term ϵ_h is single-step-error associated with h during refinement.

Enhancement of A Predictive Model In Theorem 1, $h^e(\pi_i)$ is an online approximated version of $h^*(\pi_i)$. This approximation uses the parent and the best child relationships to measure the step-error of each refinement in the path and correct it for the evaluations of further refinements. Using Theorem 1 and Figure 1, we calculate the average-step-error associated with h, denoted by ϵ_h^{avg}, along the path from π_i to π_{sol} as,

$$\epsilon_h^{avg} \;= \sum_{\pi' \textit{ from } \pi_i \rightsquigarrow \pi_{sol}} \epsilon_{h(\pi')} \Big/ h^e(\pi_i) \tag{4}$$

Rewriting Eq. (4),

$$\sum_{\pi' \textit{ from } \pi_i \rightsquigarrow \pi_{sol}} \epsilon_{h(\pi')} = \epsilon_h^{avg} \times h^e(\pi_i) \tag{5}$$

Using Eq. (5), Eq. (3) simplifies to,

$$h^e(\pi_i) \;=\; h(\pi_i) + \epsilon_h^{avg} \times h^e(\pi_i) \tag{6}$$

Further simplification of Eq. (6) yields,

$$h^e(\pi_i) = \; h(\pi_i) \Big/ (1 - \epsilon_h^{avg}) \tag{7}$$

Another possible expansion, using infinite geometric progression, of Eq. (8) would be,

$$h^e(\pi_i) = \; h(\pi_i) \times \sum_{i=0}^{\infty} (\epsilon_h^{avg})^i \tag{8}$$

We use RegPOCL to test the effectiveness of $h^e(\pi_i)$ in the POCL framework, where it selects the best partial plan.

Experiment Design

In this section we describe the evaluation phase settings. This includes (*i*) the heuristics selected as features, and (*ii*) the domains selected.

Feature Heuristics for Learning

The features used for learning are non-temporal heuristics from the literature of POCL planning. Considering the applicability of some of the POCL heuristics in the literature (Younes and Simmons 2003; Nguyen and Kambhampati 2001), we select six different heuristic functions. Some of these heuristics are informed but their informativeness varies over different planning domains. Our aim is to learn a more informed combination from these individual heuristics. The six heuristics are,

G Value ($h_{g\text{-}val}$) This returns the number of actions in a selected partial plan π not counting the two dummy actions (a_0 and a_∞). It signifies how far the search has progressed from the starting state.

Number of Open Conditions (h_{OC}) This is total number of unsupported causal links present in a partial plan, $h_{OC}(\pi) = |\text{OC}|$ (Nguyen and Kambhampati 2001).

Additive Heuristic (h_{add}) The additive heuristic h_{add} (Haslum and Geffner 2000), adds up the steps required by each individual open goal. Younes and Simmons (2003) use an adapted version of additive heuristic in POCL planning for the first time.

Additive Heuristic with Effort ($h_{add,w}$) The estimate is similar to h_{add} but it considers the cost of an action as the number of preconditions of that action, plus the linking cost **1** if the action supports any unsupported causal link (Younes and Simmons 2003). We call it $h_{add,w}$ as its notation is not used earlier. Here, w signifies the extra work required.

Accounting for Positive Interaction (h^r_{add}) This returns an estimate which takes into account the positive interactions between subgoals while ignoring the negative interactions. This is represented as h^r_{add} that is a variant of h_{add} (Younes and Simmons 2003).

Accounting for Positive Interaction with Effort ($h^r_{add,w}$) This is similar to the above heuristic which considers the total effort required (Younes and Simmons 2003). A standard notation of this heuristic is also not used in the literature.

Domains Selected

We consider the following domains: Logistics and Gripper from IPC 1, Logistics and Elevator from IPC 2, Rovers and Zenotravel from IPC 3, and Rovers from IPC 5. In our experiments we do not consider other domains from these competitions because either the state-of-the-art heuristics are not able to create enough training instances for learning, or RegPOCL does not support the domain definition language features. IPC 4 domains are not selected since the planner is not able to generate enough instances to initiate offline learning. The domains from IPC 6 and later are not supported by RegPOCL because the representations use action costs, fluents, and hard and soft constraints. Some of them can be included by some preprocessing like removal of actions with cost from the domain description files.

For each selected domain, we consider problems that are represented in STRIPS style. We select small sized problems for learning and test the learned predictive models over large sized problems in the same domain. We have a total of 109 small sized problems from the selected domains. The last four feature heuristics from the previous subsection have been used for calculating targets in each domain. This means that we generate four different datasets in each selected domain from which best two are selected. We choose satisficing track problems for generating training instances. For the training set preparation, we fix a time limit of 3 minutes and an upper limit of 500,000 on the node generation. We generate a few thousand training instances except for the Zenotravel domain where the total instances are 950. To avoid overfitting, we pick training instances between 250 to 350 from the larger training sets.

Selected Learning Approaches

In this section, we discuss in brief a procedure for feature selection in each dataset for training regression models, and different regression techniques with their references.

Feature Selection

In general, the training sets contain irrelevant or redundant attributes (out of the six selected heuristics). To reduce the training effort and increase the efficiency of our planner, we discard them from the training set. The planner is bound to calculate all the selected features at each stage of refinement. The correlation based feature selection technique (Hall 2000) is used to find the correlated features.

Regression Techniques

We use the following regression techniques to learn predictive models. These techniques have been applied in planning for learning in recent years (Samadi, Felner, and Schaeffer 2008; Thayer, Dionne, and Ruml 2011; Virseda, Borrajo, and Alcázar 2013).

Linear Regression (LR) The regression model learns a linear function that minimizes the sum of squared error over the training instances (Bishop 2006).

M5P M5P gives more flexibility than LR due to its nature of capturing non linear relationships. M5P technique learns a regression tree (Quinlan and others 1992) that approximates the class value.

M5Rules Similar to M5P but generates rules instead of modeling regression trees (Quinlan and others 1992).

Least Median Squared (LMS) LMS is similar to LR with median squared error. Functions are generated from subsamples of data with least squared error function. Usually a model with lowest median squared error is selected (Rousseeuw and Leroy 2005).

Multilayer Perceptron (MLP) MLP can learn more complex relationships compared to the other four regression techniques (Bishop 2006).

The techniques discussed above are used to learn models through WEKA (Hall et al. 2009) using a 10-fold cross-validation in each domain. A regression technique called SVMreg that implements support vector machine for regression purposes, is not used in this work due to some technical difficulty. However, it has not much influenced planning processes in the past (Virseda, Borrajo, and Alcázar 2013).

Experimental Evaluation

We use MC-Loc and MW-Loc (Younes and Simmons 2003) as flaw selecting heuristics for refining a partial plan. They give higher preference to the local flaws present in the partial plan. We employ Greedy Best First Search algorithm for selecting the next partial plan for refinement.

Environment

We perform the experiments on Intel Core 2 Quad with 2.83 GHz 64-bit processor and 4GB of RAM. To evaluate the effectiveness of learned models and to correct the single-step-error associated with the models, a time limit of 15 minutes and a node generation limit of 1 million is used.

Evaluations

We use RegPOCL to compare the performances of the offline predictive models h^l, their corresponding enhanced models $h^{l,e}$ and the last four base features. These are also compared with some of the recent effective state-space based heuristics and approaches that are introduced later. We exclude the first two base features from comparison since they are weak heuristics and RegPOCL does not solve sufficient problems using them. However, they are useful while working jointly with other informed heuristics. Next, we discuss the observations made during the training phase.

Training Using Algorithm 1, we prepare datasets and learn different predictive models by applying the various regression techniques discussed earlier. We select each of the last four features to solve a set of problems. The target value is the plan length found by RegPOCL using the base features. The dataset preparation phase took less than two

Domain	#	POCL Heuristics								Evaluation using Fast Downward (FD)								LAMA
		State-of-the-art				Via learning approaches				Lazy				Eager				
		h_{add}	$h_{add,w}$	h^r_{add}	$h^r_{add,w}$	h^l_{add}	$h^{l,e}_{add}$	$h^l_{add,w}$	$h^{l,e}_{add,w}$	FF	CEA	LM-Cut	MHFS	FF	CEA	LM-Cut	MHFS	
Gripper-1	20	16	**20**	1	1	**20**	**20**	**20**	**20**	**20**	**20**	**20**	**20**	**20**	**20**	**20**	**20**	**20**
Rovers-3	20	19	19	**20**	**20**	**20**	**20**	**20**	**20**	18	17	14	**20**	19	18	15	**20**	**20**
Rovers-5	40	28	31	32	36	$^{+}39$	$^{+}39$	$^{+}36$	$^{+}39$	22	25	15	28	24	29	17	30	**40**
		h_{add}	$h_{add,w}$	h^r_{add}	$h^r_{add,w}$	$h^{r,l}_{add}$	$h^{r,l,e}_{add}$	$h^{r,l}_{add,w}$	$h^{r,l,e}_{add,w}$	FF	CEA	LM-Cut	MHFS	FF	CEA	LM-Cut	MHFS	
Logistics-1	35	25	1	32	28	$^{+}32$	$^{+}32$	22	$^{+}32$	25	34	21	28	28	34	16	25	**35**
Elevator-2	150	148	14	**150**	58	**150**	**150**	**150**	**150**	**150**	**150**	**150**	**150**	**150**	**150**	**150**	**150**	**150**
Logistics-2	40	36	12	36	34	**40**	**40**	**40**	**40**	36	36	35	36	36	36	36	36	**40**
Zenotravel-3	20	5	4	9	10	$^{+}16$	$^{+}16$	$^{+}16$	$^{+}16$	**20**	**20**	17	**20**	**20**	**20**	16	**20**	**20**

Table 1: Number of solved problems using each heuristic. # is the number of problems in each domain. State-of-the-art POCL heuristics are compared with the learned ones in the left half. The state-based heuristics FF, CEA, and LM-Cut and their combination using MHFS strategy are also compared with POCL heuristics. The last column captures the performance of LAMA11. Best results are shown in **bold**, and a number with "+" mark (*e.g.*$^{+}$36) shows competitive performance by the learned models and their enhancements over each base heuristic. Similar representations are followed in other tables.

hours on an average in each domain, for each of the four features. Once we have enough instances, we begin the training process. In each domain, the best two datasets out of four with sufficient training points is selected for model training by visualizing the distribution of training points using WEKA. Note that, in Figure 1, different heuristics for calculating target values will prefer different paths. Therefore, the four base features will generate four different datasets. The attribute selection strategy allows us to select a subset of attributes in the training set by removing correlated and redundant features. We learn a total of 70 (7 domains $\times$ 2 datasets $\times$ 5 regression techniques) regression models. The training phase took 20ms (milliseconds) using LR, 270ms using LMS, 600ms using MLP, 82ms using M5Rule, and 58ms using M5P on an average per model. All the learned models have high accuracy but LR is the fastest, followed by M5P and M5Rule. Next, we test these models on different benchmarks.

Testing We test the effectiveness of our approaches by selecting partial plans (π) for refinement using RegPOCL. We assume that an informed heuristic leads to minimal possible refinements needed for π. Next, for the comparison we compute score[1] as in IPC for satisficing track problems. The better score on each standard signifies the better performance. We compare the performance of the learned models with the selected base features h_{add}, $h_{add,w}$, h^r_{add}, and $h^r_{add,w}$. The comparison is done on the basis of (*i.*) the number of solved problems, and (*ii.*) the score obtained on plan quality, execution time, nodes (partial plan) visited, and makespan quality.

For example, the offline learned model $h^{r,l}_{add}$ is learned on a dataset prepared using h^r_{add}. In other words, RegPOCL uses h^r_{add} for calculating the target values in the dataset. $h^{r,l}_{add}$ can be enhanced to $h^{r,l,e}_{add}$ using the online heuristic adjustment approach which is expected to be more informed than $h^{r,l}_{add}$. It is similar for other learned heuristics. These models are applied in the POCL framework for selecting the most suitable partial plan, followed by the heuristic MW-Loc (Younes and Simmons 2003) for selecting the most adverse flaw in it .

We also compare our approaches with state-space based approaches on the basis of the number of problems solved, and score obtained on plan quality and total execution time. We select fast forward heuristic (FF) (Hoffmann and Nebel 2001), context-enhanced additive heuristic (CEA) (Helmert and Geffner 2008), and landmark-cut heuristic (LM-Cut) (Helmert and Domshlak 2011). We also use these heuristics together by applying them in multi-heuristic first solution strategy (MHFS) (Röger and Helmert 2010). In general, the strategy performs better with alternating usage of different heuristics instead of combining them. We also compare the effectiveness of our techniques with LAMA11 (Richter, Westphal, and Helmert 2011); the winner of IPC-2011 in the sequential satisficing track. LAMA11 applies FF and LM-Count (Richter, Helmert, and Westphal 2008) heuristics together using multi-queue search. We set a 20 minute time limit while evaluating LAMA11 over these domains, since it has an internal time limit of 5 minutes for the invariant synthesis part of translator. All the state-based approaches are evaluated using Greedy Best First Search algorithm in the fast downward planning system (Helmert 2006). We use "eager" and "greedy" types of evaluations with no preferred operators. The regression models selected in the chosen domains are trained using, (*i*) M5P in Gripper-1 and Elevator-2, (*ii*) LR in Rovers-3, Rovers-5, Logistics-2, and Zenotravel-3, and (*iii*) M5Rule in Logistics-1.

In Table 1, we compare (*i*) the base features, (*ii*) offline learned models and their enhanced versions (*iii*) state-space based heuristics FF, CEA, and LM-Cut, and (*iv*) the strategies used in MHFS, and LAMA11, on the basis of number of problems solved. In this table, RegPOCL solves equal number of problems as LAMA11 in Gripper-1, Rovers-3, Elevator-2, and Logistics-2 using each of learned heuristics and their enhanced versions. The base features have performed well in some domains but are not consistent overall. In Rovers-5, our approaches solved 1 problem less than LAMA11, but they beat other state-space based competitors comprehensively. Also, each learned model has performed better than all the base features. For Logistics-2, we are competitive with LAMA11 and solve at least 4 more problems than other good heuristics like CEA and LM-Cut. In Zenotravel-3, RegPOCL solved 6 problems more by applying our approaches but loses to the state-based competitors. Our second approach improves the performance of the

[1] https://helios.hud.ac.uk/scommv/IPC-14/

Domain	POCL Heuristics								Evaluation using Fast Downward (FD)								LAMA
	State-of-the-art				Via learning approaches				Lazy				Eager				
	h_{add}	$h_{add,w}$	h^r_{add}	$h^r_{add,w}$	h^l_{add}	$h^{l,e}_{add}$	$h^l_{add,w}$	$h^{l,e}_{add,w}$	FF	CEA	LM-Cut	MHFS	FF	CEA	LM-Cut	MHFS	
Gripper-1	16.0	14.9	0.7	0.7	**20.0**	**20.0**	**20.0**	**20.0**	15.45	14.9	15.5	15.5	15.5	14.9	15.5	15.5	**20.0**
Rovers-3	17.3	16.2	18.1	16.9	17.8	+18.6	17.8	+18.6	17.1	15.9	12.5	18.8	18.1	16.9	13.9	19.0	**19.8**
Rovers-5	25.7	26.3	28.0	30.2	+33.1	+36.3	30.1	+33.6	20.8	23.4	13.5	26.0	22.7	27.3	15.7	28.3	**39.8**
	h_{add}	$h_{add,w}$	h^r_{add}	$h^r_{add,w}$	$h^{r,l}_{add}$	$h^{r,l,e}_{add}$	$h^{r,l}_{add,w}$	$h^{r,l,e}_{add,w}$	FF	CEA	LM-Cut	MHFS	FF	CEA	LM-Cut	MHFS	
Logistics-1	24.3	0.8	31.2	26.1	+31.2	+31.2	20.9	**32.0**	23.7	30.3	19.7	27.4	27.0	30.8	15.1	24.6	30.8
Elevator-2	142.8	11.8	144.9	49.7	142.7	+144.9	142.7	+144.9	117.1	112.0	117.1	116.9	144.1	136.0	144.2	141.0	**148.2**
Logistics-2	33.8	10.7	34.9	30.7	**38.1**	+36.3	**38.1**	+36.3	33.72	29.5	32.4	33.3	34.1	30	33.7	35.6	37.8
Zenotravel-3	4.6	3.7	8.5	8.8	+13.5	+13.7	+13.5	+13.7	17.3	17.7	14.7	18.6	18.1	14.3	16	18.8	**19.2**
Time Score	204.4	76.5	197.7	148.6	+272.9	+260.8	+258.6	+264.3	255.6	**280.0**	228.0	233.8	259.6	271.3	223.0	222.5	137

Table 2: Scores on plan quality and overall time. We compare state-of-the-art POCL heuristics with learned ones. The effectiveness of the POCL heuristics is compared with some latest state based approaches. The last row demonstrates the overall time score of each heuristic. The numerical representations and column details are similar to Table 1.

Domain	State-of-the-art				Via learning approaches			
	h_{add}	$h_{add,w}$	h^r_{add}	$h^r_{add,w}$	h^l_{add}	$h^{l,e}_{add}$	$h^l_{add,w}$	$h^{l,e}_{add,w}$
Gripper-1	7.0	14.0	0.0	0.0	+18.0	**19.7**	+18.0	**19.7**
Rovers-3	8.6	8.3	12.6	**19.0**	13.3	13.6	13.3	13.6
Rovers-5	8.2	14.7	13.5	30.9	**31.4**	**31.4**	24.7	**31.4**
	h_{add}	$h_{add,w}$	h^r_{add}	$h^r_{add,w}$	$h^{r,l}_{add}$	$h^{r,l,e}_{add}$	$h^{r,l}_{add,w}$	$h^{r,l,e}_{add,w}$
Logistics-1	5.1	0.6	16.5	21.6	**29.9**	**29.9**	20.7	**29.9**
Elevator-2	41.4	8.2	58.9	39.7	**145.0**	+135.0	**145.0**	+135.0
Logistics-2	24.5	6.5	23.4	26.7	+33.4	**35.3**	+33.4	**35.3**
Zenotravel-3	0.0	1.0	2.4	7.7	7.4	**11.07**	7.4	**11.07**

Table 3: Results of refinement score (the number nodes visited) of RegPOCL using each heuristic. We compare state-of-the-art heuristics with the learned models (h^l) and their further enhancements ($h^{l,e}$). The best results are in **bold**.

Domain	h_{add}	$h_{add,w}$	h^r_{add}	$h^r_{add,w}$	h^l_{add}	$h^{l,e}_{add}$	$h^l_{add,w}$	$h^{l,e}_{add,w}$
Gripper-1	16.0	9.8	0.5	0.5	**20.0**	**20.0**	**20.0**	**20.0**
Rovers-3	**17.9**	15.5	17.8	15.8	17.4	17.7	17.4	17.7
Rovers-5	26.1	24.2	28.4	26.5	**30.5**	+29.8	+27.6	**30.5**
	h_{add}	$h_{add,w}$	h^r_{add}	$h^r_{add,w}$	$h^{r,l}_{add}$	$h^{r,l,e}_{add}$	$h^{r,l}_{add,w}$	$h^{r,l,e}_{add,w}$
Logistics-1	21.8	0.9	25.2	20.8	+30.4	+30.4	17.6	**30.5**
Elevator-2	147.1	11.7	**149.1**	50.5	**149.1**	146.8	**149.1**	146.8
Logistics-2	26.3	7.0	33.0	19.3	+33.5	**35.4**	+33.4	**35.4**
Zenotravel-3	3.7	2.7	7.5	7.7	+12.1	**13.5**	+12.1	**13.5**

Table 4: Makespan quality score of the heuristics. Columns are identical to corresponding ones in Table 3.

learned models in Rovers-5 by solving 3 more problems, and in Logistics-1 where it solves 10 more problems. This approach could not increase the coverage in other domains. LAMA11 wins on the basis of the total number of problems solved in each domain.

In Table 2, we compare the score obtained on plan quality by each of the base features, learned models with their corresponding enhancements, and state-space based heuristics and techniques. LAMA11 is an anytime planner which gives solution close to the optimal but takes more time to compute shorter plans. For Gripper-1 and Logistics-2, RegPOCL using the learned heuristics solves equal number of problems but finds shorter plans compared to LAMA11. In Logistics-1, RegPOCL lost to LAMA11 by 3 problems in the number of problems solved, but obtained a higher score in the other problems it solved as it produces shorter plans using $h^{r,l,e}_{add,w}$, $h^{r,l}_{add}$, and $h^{r,l,e}_{add}$. The effectiveness of $h^{r,l,e}_{add,w}$ wins in this domain with best plan quality. $h^{r,l,e}_{add,w}$ also increases the plan quality score from 20.9 that is obtained by $h^{r,l}_{add,w}$ to 32.0 using our online error correcting strategy. However, in Logistics-2, the performance has decreased after further enhancements of $h^{r,l}_{add}$ and $h^{r,l}_{add,w}$.

In general, the performance of RegPOCL using either of the offline learned models and their enhancements is often better than that using the base features. In most of the cases, the online error adjustment approach has further enhanced the performance of these learned models. The last row of Table 2 gives the score obtained on total time taken by the process. If a planner takes less than 1 second for solving a problem then it gets full score. On the total time score the winner is CEA with "lazy" evaluation. The learned models and their enhanced versions have obtained better scores than other competitors except CEA. These scores are very close to the winning score and almost twice that of LAMA11.

In Table 3, we compare the score obtained on the number of nodes RegPOCL visits for solving the planning problems. This is obtained for the base features, the learned heuristics, and their enhanced versions. The models obtained after offline learning are more informed toward goals and refines fewer partial plans. The score obtained by the learned models is further increased by a good factor in Zenotravel-3 using the error correcting approach. For Elevator-2, the error correcting approach has shown some negative effect which continues in Table 4 too. In Table 4, we demonstrate the score obtained on the makespan quality. Higher score signifies smaller makespan and more flexibility in the generated solution plan. In Elevator-2 and Rovers-5, the scores of h^l_{add} and $h^{r,l}_{add}$ have decreased due to the negative effects of our error adjustment approach, while the score obtained by $h^{r,l,e}_{add,w}$ is almost 1.5 times the score of $h^{r,l}_{add,w}$ in Logistics-1. In general, the offline learned models have generated more flexible plans with shorter makespan than the base features. These qualities are further improved using the enhanced versions of these models.

Discussion

We have already discussed the advantages of our approaches but they also have limitations. In our offline approach, we are bound to do some poor generalization while learning heuristics. Current literature supports the idea of selecting a large feature set for more accurate learning (Roberts et al. 2008). Accuracy can also be improved using an empirical

performance model of all components of a portfolio to decide which component to pick next (Fawcett et al. 2014). In our work, a large feature set may have some drawbacks. For example, computing the features at each refinement step during the planning process is computationally expensive.

The online error adjustment approach could also perform poorly in certain domains. In Figure 1, if the orientation of objects in a domain is such that $h(\pi_{i+1})$ is larger than $h(\pi_i)$ then $\epsilon_{h(\pi_i)}$ may not be accurate. The inaccuracy in $\epsilon_{h(\pi_i)}$ is compounded if the above condition holds at the beginning of the planning process. This results in an inaccurate ϵ_h^{avg} value, leading to wrong selection of the partial plan to refine next. Consequently, the planner may end up finding longer and less flexible plans. Another limitation is that a refinement may change the existing priorities of partial plans in the set due to the single-step-error adjustment. Considering the time factor, we avoid changing the decided priorities of those partial plans. This may also lead to inaccuracy in ϵ_h^{avg}.

Our approaches do not utilize the advantage of strategies like alternation queue, and candidate selection using concept of pareto optimality (Röger and Helmert 2010). Recently, the planning community has tried coming up with effective portfolios of heuristics or planners. The techniques of generating good portfolios are not new to theoretical machine learning. A follow up work done in the past is combining multiple heuristics online (Streeter, Golovin, and Smith 2007). One could form a portfolio of different algorithms to reduce the total makespan for a set of jobs to solve (Streeter and Smith 2008). The authors provide a bound on the performance of the portfolio approaches. For example, an execution of a greedy schedule of algorithms cannot exceed four times the optimal schedule.

In planning, a sequential portfolio of planners or heuristics aims to optimize the performance metrics. In general, such configurations automatically generate sequential orderings of best planning algorithms. In the portfolio the participants are allotted some timestamp to participate in solving problems in the ordering. A similar approach is used in (Seipp et al. 2015). The authors outline their procedure for optimal and satisficing planning. The procedure used in this work starts with a set of planning algorithms and a time bound. It uses another procedure *OPTIMIZE* that focuses on the marginal improvements of the performance. Here, the quality of the portfolio is bounded by $(1 - (1/\mathrm{e})) \times \mathrm{OPT}$, and the running time cannot exceed 4 OPT. The components can be allowed to act in a round-robin fashion (Gerevini, Saetti, and Vallati 2014).

The state-of-the-art planners exhibit variations in their runtime for a given problem instance, so no planner always dominates over others. A good approach would be to select a planner for a given instance by looking at its processing time. This is done by building an empirical performance model (EPM) for each planner. EPM is derived from sets of planning problems and performance observation. It predicts whether the planner could solve a given instance (Fawcett et al. 2014). The authors consider a large set of instance features and show that the runtime predictor is often superior to the individual planners. Performance wise sorting of components in a portfolio is also possible (Núñez, Borrajo, and López 2015). The portfolio is sorted such that the probability of the performance of that portfolio is maximum at any time. Experiments show that performance of a greedy strategy can be enhanced to near optimal over time.

The last two paragraphs cover recent literature in brief which explain previous strategies of combining different base methods. The literature shows that they have performed well over different benchmarks. Our current settings do not capture any such ideas for combining different components of heuristics. A direct comparison with any of the above mentioned works is therefore out of scope for our current work. This is because, we are more concerned about working with unit cost based POCL heuristics in isolation. On the other hand, we suspect that many of these strategies, in some adapted form, would likely be beneficial in the POCL framework.

Summary and Future Work

We demonstrate the use of different regression models to combine different heuristic values to arrive at consistently better estimates over a range of planning domains and problems. We extend some recent attempts to learn combinations of heuristic functions in state-space based planning to POCL planning. We also show that the learned models can be further enhanced by an online error correction approach.

In future we intend to explore online learning further, and continue our experiments with combining heuristic functions. We also aim to explore the use of an optimizing planner in tandem with bootstrapping methods. Apart from these, we will be giving a complete generalization of our current learning approaches for temporal planning and planning with deadlines.

References

[Arfaee, Zilles, and Holte 2011] Arfaee, S. J.; Zilles, S.; and Holte, R. C. 2011. Learning heuristic functions for large state spaces. *Artificial Intelligence* 175(16-17):2075–2098.

[Benton, Coles, and Coles 2012] Benton, J.; Coles, A. J.; and Coles, A. 2012. Temporal planning with preferences and time-dependent continuous costs. In *Proc. ICAPS 2012*, 2–10.

[Bishop 2006] Bishop, C. M. 2006. *Pattern Recognition and Machine Learning*, volume 1. Springer, Aug. 2006.

[Coles et al. 2010] Coles, A. J.; Coles, A.; Fox, M.; and Long, D. 2010. Forward-chaining partial-order planning. In *Proc. ICAPS*.

[Domshlak, Karpas, and Markovitch 2010] Domshlak, C.; Karpas, E.; and Markovitch, S. 2010. To max or not to max: Online learning for speeding up optimal planning. In *Proc. AAAI 2010*, 1071–1076.

[Fawcett et al. 2014] Fawcett, C.; Vallati, M.; Hutter, F.; Hoffmann, J.; Hoos, H. H.; and Leyton-Brown, K. 2014. Improved features for runtime prediction of domain-independent planners. In *Proc. ICAPS 2014*, 355–359.

[Gerevini, Saetti, and Vallati 2014] Gerevini, A.; Saetti, A.; and Vallati, M. 2014. Planning through automatic portfolio configuration: The PbP approach. *JAIR 2014*. 639–696.

[Hall et al. 2009] Hall, M.; Frank, E.; Holmes, G.; Pfahringer, B.; Reutemann, P.; and Witten, I. H. 2009. The WEKA data mining software: An update. In *SIGKDD Explorations 2009*, volume 11.

[Hall 2000] Hall, M. A. 2000. Correlation-based feature selection for discrete and numeric class machine learning. In *Proc. ICML*.

[Haslum and Geffner 2000] Haslum, P., and Geffner, H. 2000. Admissible heuristics for optimal planning. In *Proc. AIPS 2000*, 140–149.

[Haslum 2012] Haslum, P. 2012. Incremental lower bounds for additive cost planning problems. In *Proc. ICAPS 2012*, 74–82.

[Helmert and Domshlak 2011] Helmert, M., and Domshlak, C. 2011. LM-Cut: Optimal planning with the landmark-cut heuristic. In *Seventh IPC*.

[Helmert and Geffner 2008] Helmert, M., and Geffner, H. 2008. Unifying the causal graph and additive heuristics. In *Proc. ICAPS 2008*.

[Helmert, Röger, and Karpas 2011] Helmert, M.; Röger, G.; and Karpas, E. 2011. Fast downward stone soup: A baseline for building planner portfolios. In *Proc. ICAPS PAL 2011*, 28–35. Citeseer.

[Helmert 2006] Helmert, M. 2006. The fast downward planning system. *J. Artif. Intell. Res. (JAIR) 2006,* 26:191–246.

[Hoffmann and Nebel 2001] Hoffmann, J., and Nebel, B. 2001. The FF planning system: Fast plan generation through heuristic search. *JAIR 2001*.

[Kvarnström 2011] Kvarnström, J. 2011. Planning for loosely coupled agents using partial order forward-chaining. In *Proc. ICAPS 2011*.

[Nguyen and Kambhampati 2001] Nguyen, X., and Kambhampati, S. 2001. Reviving partial order planning. In *Proc. IJCAI 2001*, 459–466.

[Núñez, Borrajo, and López 2015] Núñez, S.; Borrajo, D.; and López, C. L. 2015. Sorting sequential portfolios in automated planning. In *Proc. IJCAI 2015*, 1638–1644. AAAI Press.

[Quinlan and others 1992] Quinlan, J. R., et al. 1992. Learning with continuous classes. In *Proc. Australian Joint Conference on Artificial Intelligence*, volume 92, 343–348.

[Richter and Westphal 2010] Richter, S., and Westphal, M. 2010. The LAMA planner: Guiding cost-based anytime planning with landmarks. *JAIR 2010,* 39:127–177.

[Richter, Helmert, and Westphal 2008] Richter, S.; Helmert, M.; and Westphal, M. 2008. Landmarks revisited. In *Proc. AAAI 2008*, volume 8, 975–982.

[Richter, Westphal, and Helmert 2011] Richter, S.; Westphal, M.; and Helmert, M. 2011. LAMA 2008 and 2011. In *Seventh IPC 2011*, 117–124.

[Roberts et al. 2008] Roberts, M.; Howe, A. E.; Wilson, B.; and desJardins, M. 2008. What makes planners predictable?. In *Proc. ICAPS 2008*, 288–295. AAAI.

[Röger and Helmert 2010] Röger, G., and Helmert, M. 2010. The more, the merrier: Combining heuristic estimators for satisficing planning. In *Proc. ICAPS 2010*, 246–249.

[Röger, Pommerening, and Seipp 2014] Röger, G.; Pommerening, F.; and Seipp, J. 2014. Fast downward stone soup 2014. In *Eighth IPC 2014*.

[Rousseeuw and Leroy 2005] Rousseeuw, P. J., and Leroy, A. M. 2005. *Robust regression and outlier detection*, volume 589. John Wiley & Sons.

[Samadi, Felner, and Schaeffer 2008] Samadi, M.; Felner, A.; and Schaeffer, J. 2008. Learning from multiple heuristics. In *Proc. AAAI 2008*, 357–362.

[Sapena, Onaindıa, and Torreno 2014] Sapena, O.; Onaindıa, E.; and Torreno, A. 2014. Combining heuristics to accelerate forward partial-order planning. In *Proc ICAPS COPLAS 2014*.

[Seipp et al. 2015] Seipp, J.; Sievers, S.; Helmert, M.; and Hutter, F. 2015. Automatic configuration of sequential planning portfolios. In *Proc. AAAI 2015*, 3364–3370.

[Shekhar and Khemani 2015] Shekhar, S., and Khemani, D. 2015. Improving heuristics on-the-fly for effective search in plan space. In *KI 2015: Advances in Artificial Intelligence - 38th Annual German Conference on AI*, 302–308.

[Streeter and Smith 2008] Streeter, M. J., and Smith, S. F. 2008. New techniques for algorithm portfolio design. In *Proc. UAI 2008*, 519–527.

[Streeter, Golovin, and Smith 2007] Streeter, M. J.; Golovin, D.; and Smith, S. F. 2007. Combining multiple heuristics online. In *Proc. AAAI 2007*.

[Sutton 1988] Sutton, R. S. 1988. Learning to predict by the methods of temporal differences. *Machine Learning,* 3:9–44.

[Thayer, Dionne, and Ruml 2011] Thayer, J. T.; Dionne, A. J.; and Ruml, W. 2011. Learning inadmissible heuristics during search. In *Proc. ICAPS 2011*.

[Virseda, Borrajo, and Alcázar 2013] Virseda, J.; Borrajo, D.; and Alcázar, V. 2013. Learning heuristic functions for cost-based planning. In *Proc. ICAPS PAL 2013*. Citeseer.

[Younes and Simmons 2003] Younes, H. L. S., and Simmons, R. G. 2003. VHPOP: Versatile Heuristic Partial Order Planner. *JAIR 2003,* 20:405–430.

Appendix

Proof of Theorem 1

Theorem 1. *For a given learned predictive model (h) and partial plan (π_i) in Figure 1 which leads to the solution plan (π_{sol}) after certain refinement steps, the enhanced version of the predictive model (h^e) is,*

$$h^e(\pi_i) \;=\; h(\pi_i) \;+\; \sum_{\substack{\pi' \text{ from} \\ \pi_i \rightsquigarrow \pi_{sol}}} \epsilon_{h(\pi')} \qquad (9)$$

where $\pi_i \rightsquigarrow \pi_{sol}$ is a path in Figure 1 which captures each partial plan (π') along the path between π_i and π_{sol}. The path includes π_i and excludes the π_{sol}. The term ϵ_h is single-step-error associated with h during refinement.

Proof. We use the principle of mathematical induction to prove this theorem.

Basis : We assume that π_i needs only one refinement to become π_{sol} that would also be π_i's best child. Here, the best child always keeps the lowest estimate of requirement of new actions for its refinement among its siblings. One possible refinement of π_i is, $\pi_i \xrightarrow{R_i} \pi_{sol}$. Using Eq. (9), we say,

$$h^e(\pi_i) \quad = \quad h(\pi_i) \; + \; \epsilon_{h(\pi_i)} \tag{10}$$

The term $\epsilon_{h(\pi_i)}$ is the single-step-error associated with *h* that estimates the total effort required for π_i to refine it completely. For unit cost refinements (cost(R_i) = 1), $\epsilon_{h(\pi_i)}$ is computed as,

$$\epsilon_{h(\pi_i)} = (cost(R_i) + h(\pi_{i+1})) - h(\pi_i) \tag{11}$$

Here, the partial plan π_{i+1} is also the π_{sol}, therefore h(π_{i+1}) = 0. By using Eq. (10) and Eq. (11) together, we get $h^e(\pi_i)$ = 1. Therefore, the base step holds.

In the base case, we assume that after refinement step R_i, to be the best child, there is no unsupported causal link present in π_{i+1} and a threat (if any) will be resolved by the planner immediately to make it a solution plan. If there is an unsupported causal link then there must be an existing action in π_{i+1} to support it. In this case, the estimate of requirement of new actions is still be 0.

Hypothesis : We select an arbitrary partial plan π_{i+1} and assume that Eq. (9) holds for it.

Proof Step : Here, we show that Eq. (9) holds for π_i too.

$$\begin{aligned}
h^e(\pi_i) &= cost(R_i) + h^e(\pi_{i+1}) \\
&= cost(R_i) + h(\pi_{i+1}) + \sum_{\pi' \textit{ from } \pi_{i+1} \rightsquigarrow \pi_{sol}} \epsilon_{h(\pi')} \\
&\quad \text{By the induction hypothesis.} \\
&= h(\pi_i) + \epsilon_{h(\pi_i)} + \sum_{\pi' \textit{ from } \pi_{i+1} \rightsquigarrow \pi_{sol}} \epsilon_{h(\pi')} \text{ by Eq. (11).} \\
&= h(\pi_i) + \sum_{\pi' \textit{ from } \pi_i \rightsquigarrow \pi_{sol}} \epsilon_{h(\pi')}
\end{aligned}$$

Therefore, the the relationship holds for the parent partial plan π_i as well. Thus, by induction for all partial plans π, our assumption is correct. ∎